Terms used **multi dimensional** or **multidimensional** or **cube** or **hypercube** and **index sentence deviation** or **divergence** or **variation** or **divergency**

Found **65** of **177,263**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 65 Result page: **1** 2 3 4 next

Relevance scale

1 Achieving scalability in OLAP materialized view selection

Thomas P. Nadeau, Toby J. Teorey
November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP**
Publisher: ACM Press

Full text available: pdf(347.21 KB) Additional Information: full citation, abstract, references, citings, index terms

The goal of on-line analytical processing (OLAP) is to quickly answer queries from large amounts of data residing in a data warehouse. Materialized view selection is an optimization problem encountered in OLAP systems. Published work on the problem of materialized view selection presents solutions scalable in the number of possible views. However, the number of possible views is exponential relative to the number of database dimensions. A truly scalable solution must be polynomial time relative ...

Keywords: OLAP, OLAP performance, data warehouse, materialized views, view selection

2 Generalized multidimensional data mapping and query processing

Rui Zhang, Panos Kalnis, Beng Chin Ooi, Kian-Lee Tan
September 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 3
Publisher: ACM Press

Full text available: pdf(689.08 KB) Additional Information: full citation, abstract, references, index terms

Multidimensional data points can be mapped to one-dimensional space to exploit single dimensional indexing structures such as the B+-tree. In this article we present a Generalized structure for data Mapping and query Processing (GiMP), which supports extensible mapping methods and query processing. GiMP can be easily customized to behave like many competent indexing mechanisms for multi-dimensional indexing, such as the UB-Tree, the Pyramid technique, the iMinMax, and the iDistan ...

Keywords: Indexing, data mapping, efficiency

3 Computational investigations of low-discrepancy sequences

Ladislav Kocis, William J. Whiten

June 1997 **ACM Transactions on Mathematical Software (TOMS)**, Volume 23 Issue 2
Publisher: ACM Press

Full text available: pdf(295.58 KB) Additional Information: full citation, abstract, references, citings, index terms

The Halton, Sobol, and Faure sequences and the Braaten-Weller construction of the generalized Halton sequence are studied in order to assess their applicability for the quasi Monte Carlo integration with large number of variates. A modification of the Halton sequence (the Halton sequence leaped) and a new construction of the generalized Halton sequence are suggested for unrestricted number of dimensions and are shown to improve considerably on the original Halton sequence. Problems associat ...

Keywords: Faure sequence, Halton sequence, Monte Carlo and quasi Monte Carlo integration, Sobol sequence, discrepancy, error of numerical integration, generalized Halton sequence, low-discrepancy sequences

4 STHoles: a multidimensional workload-aware histogram

Nicolas Bruno, Surajit Chaudhuri, Luis Gravano
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2
Publisher: ACM Press

Full text available: pdf(429.21 KB) Additional Information: full citation, abstract, references, citings, index terms

Attributes of a relation are not typically independent. Multidimensional histograms can be an effective tool for accurate multiattribute query selectivity estimation. In this paper, we introduce *STHoles*, a "workload-aware" histogram that allows bucket nesting to capture data regions with reasonably uniform tuple density. *STHoles* histograms are built without examining the data sets, but rather by just analyzing query results. Buckets are allocated where needed the mos ...

5 Query evaluation techniques for large databases

Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

6 Statistical profile estimation in database systems

Michael V. Mannino, Paicheng Chu, Thomas Sager
September 1988 **ACM Computing Surveys (CSUR)**, Volume 20 Issue 3
Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index

Full text available: pdf(2.94 MB) terms

A statistical profile summarizes the instances of a database. It describes aspects such as the number of tuples, the number of values, the distribution of values, the correlation between value sets, and the distribution of tuples among secondary storage units. Estimation of database profiles is critical in the problems of query optimization, physical database design, and database performance prediction. This paper describes a model of a database of profile, relates this model to estimating ...

7 An optimal algorithm for approximate nearest neighbor searching fixed dimensions

Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, Angela Y. Wu
November 1998 **Journal of the ACM (JACM)**, Volume 45 Issue 6
Publisher: ACM Press

Full text available: pdf(287.94 KB) Additional Information: full citation, abstract, references, citings, index terms

Consider a set of S of n data points in real d-dimensional space, Rd, where distances are measured using any Minkowski metric. In nearest neighbor searching, we preprocess S into a data structure, so that given any query point q ∈ Rd, is the closest point of S to q can be reported quickly. Given any po ...

Keywords: approximation algorithms, box-decomposition trees, closet-point queries, nearest neighbor searching, post-office problem, priority search

8 Analysis methodology I: Quasi-Monte Carlo methods in cash flow testing simulations

Michael G. Hilgers
December 2000 **Proceedings of the 32nd conference on Winter simulation**
Publisher: Society for Computer Simulation International
Full text available: pdf(591.55 KB) Additional Information: full citation, abstract, references

What actuaries call *cash flow testing* is a large-scale simulation pitting a company's current policy obligation against future earnings based on interest rates. While life contingency issues associated with contract payoff are a mainstay of the actuarial sciences, modeling the random fluctuations of US Treasury rates is less studied. Furthermore, applying standard simulation techniques, such as the Monte Carlo method, to actual multi-billion dollar companies produce a simulation that can ...

9 High Dimensional Direct Rendering of Time-Varying Volumetric Data

Jonathan Woodring, Chaoli Wang, Han-Wei Shen
October 2003 **Proceedings of the 14th IEEE Visualization 2003 (VIS'03) VIS '03**
Publisher: IEEE Computer Society
Full text available: pdf(473.10 KB) Additional Information: full citation, abstract

We present an alternative method for viewing time-varying volumetric data. We consider such data as a four-dimensional data field, rather than considering space and time as separate entities. If we treat the data in this manner, we can apply high dimensional slicing and projection techniques to generate an image hyperplane. The user is provided with an intuitive user interface to specify arbitrary hyperplanes in 4D, which can be displayed with standard volume rendering techniques. From the volum ...

Keywords: time-varying data, hyperslice, hyperprojection, integration operator, transfer function, raycasting, volume rendering

10 K-d trees for semidynamic point sets

Jon Louis Bentley
May 1990 **Proceedings of the sixth annual symposium on Computational geometry**
Publisher: ACM Press

Full text available: pdf(928.78 KB) Additional Information: full citation, abstract, references, citings, index terms

A K-d tree represents a set of N points in K-dimensional space. Operations on a semidynamic tree may delete and undelete points, but may not insert new points. This paper shows that several operations that require &Ogr;(log N) expected time in general K-d trees may be performed in constant expected time in semidynamic trees. These operations include deletion, undeletion, ne ...

11 iDistance: An adaptive B^{+}-tree based indexing method for nearest neighbor search
H. V. Jagadish, Beng Chin Ooi, Kian-Lee Tan, Cui Yu, Rui Zhang
June 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 2
Publisher: ACM Press
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, index terms

In this article, we present an efficient B+-tree based indexing method, called iDistance, for K-nearest neighbor (KNN) search in a high-dimensional metric space. iDistance partitions the data based on a space- or data-partitioning strategy, and selects a reference point for each partition. The data points in each partition are transformed into a single dimensional value based on their similarity with respect to the reference point. This allows the points to be indexed using a B
~~**Keywords**: Indexing, KNN, nearest neighbor queries~~

-

12 Overlay networks, scalability and internet economics: Location based placement of whole distributed systems
David Spence, Jon Crowcroft, Steven Hand, Tim Harris
October 2005 **Proceedings of the 2005 ACM conference on Emerging network experiment and technology CoNEXT'05**
Publisher: ACM Press
Full text available: pdf(298.25 KB) Additional Information: full citation, abstract, references, index terms

The high bandwidth and low latency of the modern internet has made possible the deployment of *distributed computing platforms*. The *XenoServe* platform provides a distributed computing platform open to all and presents three major new challenges for resource discovery: Firstly, network location is key for effectively provisioning services, to mitigate against high-latency, high-load or component failure. Secondly, many services require a presence on several servers, with inter-relate ...

Keywords: location systems, peer-to-peer, resource discovery

13 A cost model for query processing in high dimensional data spaces
Christian Böhm
June 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(362.22 KB) Additional Information: full citation, abstract, references, citings, index terms, review

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research topic in multimedia databases is similarity search in large data sets. Most current approaches that address similarity search use the feature approach, which transforms important properties of the stored objects into points of a high-dimensional

space (feature vectors). Thus, similarity search is transformed ...

Keywords: cost model, multidimensional index

14 Implementing data cubes efficiently

Venky Harinarayan, Anand Rajaraman, Jeffrey D. Ullman
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

Decision support applications involve complex queries on very large databases. Since response times should be small, query optimization is critical. Users typically view the data as multidimensional data cubes. Each cell of the data cube is a view consisting of an aggregation of interest, like total sales. The values of many of these cells are dependent on the values of other cells in the data cube. A common and powerful query optimization technique is to materialize some or all of these cells r ...

15 Session 3B: Optimal online bounded space multidimensional packing

Leah Epstein, Rob van Stee
January 2004 **Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms**
Publisher: Society for Industrial and Applied Mathematics
Full text available: pdf(181.59 KB) Additional Information: full citation, abstract, references, citings

We solve an open problem in the literature by providing an online algorithm for multidimensional bin packing that uses only bounded space. To achieve this, we introduce a new technique for classifying the items to be packed. We show that our algorithm is optimal among bounded space algorithms for any dimension $d > 1$. Its asymptotic performance ratio is $(\Pi_\infty)^d$, where $\Pi_\infty \approx 1{:}691$ is the asymptotic performance ratio of the one-dimensional algorithm HARM ...

16 Clustering declustered data for efficient retrieval

Hakan Ferhatosmanoglu, Divyakant Agrawal, Amr El Abbadi
November 1999 **Proceedings of the eighth international conference on Information and knowledge management**
Publisher: ACM Press

Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references, citings, index terms

Modern databases increasingly integrate new kinds of information, such as multimedia information in the form of image, video, and audio data. Both the dimensionality and the amount of data that need to be processed is increasing rapidly, increasing the demand for the efficient retrieval of large amounts of multi-dimensional data. Declustering techniques for multi-disk architectures have been effectively used for storage. In this paper, we first establish that besides exploiting the parallel ...

17 A fiber optic hypermesh for SIMD/MIMD machines

Ted Szymanski
November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society
Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references

A fiber optic multidimensional mesh-based network for SIMD and MIMD multiprocessors is proposed. For the basic building block, a novel distributed optical switch is proposed; The

switch requires 50 % fewer lasers/receivers than previous WDM optical crossbars and uses a novel random-access scheme which supports prioritized traffic. To implement very large networks using lasers with limited tunability (or electronic crossbars of small degree) we propose arranging switches into a novel *n*-dim ...

18 From discrepancy to declustering: Near-optimal multidimensional declustering strategies for range queries

Chung-Min Chen, Christine T. Cheng
January 2004 **Journal of the ACM (JACM)**, Volume 51 Issue 1
Publisher: ACM Press
Full text available: pdf(225.33 KB) Additional Information: full citation, abstract, references, index terms

Declustering schemes allocate data blocks among multiple disks to enable parallel retrieval. Given a declustering scheme *D*, its *response time* with respect to a query *Q*, *rt* (*Q*), is defined to be the maximum number of data blocks of the query stored by the scheme in any one of the disks. If |*Q*| is the number of data blocks in *Q* and *M* is the number of disks, then *rt*(*Q*) is at least □|*Q*|/*M*□. One way to eval ...

Keywords: Declustering schemes, disk allocations, parallel database, range query

19 Searching in high-dimensional spaces: Index structures for improving the performance of multimedia databases

Christian Böhm, Stefan Berchtold, Daniel A. Keim
September 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 3
Publisher: ACM Press
Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research issue in the field of multimedia databases is the content-based retrieval of similar multimedia objects such as images, text, and videos. However, in contrast to searching data in a relational database, a content-based retrieval requires the search of similar objects as a basic functionality of the database system ...

Keywords: Index structures, indexing high-dimensional data, multimedia databases, similarity search

20 The GOLD definition language (GDL): an object oriented formal specification language for multidimensional databases

Juan Trujillo, Manuel Palomar, Jaime Gómez
March 2000 **Proceedings of the 2000 ACM symposium on Applied computing - Volume 1**
Publisher: ACM Press
Full text available: pdf(421.67 KB) Additional Information: full citation, references, citings, index terms

Keywords: OLAP, conceptual modeling, data warehouses, multidimensional databases, object-orientation

ProQuest®

Searching for *(multidimensional database* or cube* or hypercube*) and (creat* or generat*) w/20 (index*) w/20 (deviation* or variation* or diverg*)* did not find any documents. Try the following:

Revise your search below using the following tips:

- Check your spelling.
- Reduce the number of terms included in your search.
- Broaden your search by selecting other databases, removing limits, or searching "Citations and Document Text" (if available).
- Use "AND" to connect two words that don't need to be searched as a phrase.
- Connect similar terms with the "OR" operator (e.g. military OR pentagon). See Search Tips for more hints.

Basic Search

Tools: Search Tips | Browse Topics | 1 Recent Searches

(multidimensional database* or cube* or hypercube*) and (creat* or generat | Search | Clear

Database: Multiple databases... Select multiple databases

Date range: All dates

Limit results to: ☐ Full text documents only

☐ Scholarly journals, including peer-reviewed About

More Search Options

Text-only interface

ProQuest COMPANY

ProQuest®

Results

10 documents found for: *(multidimensional database* or cube* or hypercube*) and (index*) and (deviation* or variation* or diverg*)* » Refine Search | Set Up Alert

All sources | Scholarly Journals | Newspapers | Dissertations

Mark all 0 marked items: Email / Cite / Export Show only full text Sort results by: Most recent first

1. **Cubes and the individual investor**
Richard J Curcio, Joanna M Lipka, John H Thornton Jr. **Financial Services Review.** Atlanta: Summer 2004. Vol. 13, Iss. 2; p. 123 (16 pages)
Text+Graphics Page Image - PDF Abstract

2. **Shoreline response to variations in waves and water levels: An engineering scale approach**
by *Miller, Jonathan K.*, Ph.D., **University of Florida**, 2004, 175 pages; AAT 3161019
Abstract 24 Page Preview Page Image - PDF Order a copy

3. **Cluster-based storage policies in kitting area**
by *Hua, Wei*, Ph.D., **Georgia Institute of Technology**, 2001, 125 pages; AAT 3015657
Abstract 24 Page Preview Page Image - PDF Order a copy

4. **Parameter identification and sensitivity analysis of the differential model of hysteresis**
by *Zhang, Haochuan*, Ph.D., **University of California, Berkeley**, 2001, 144 pages; AAT 3044751
Abstract 24 Page Preview Page Image - PDF Order a copy

5. **House price inflation still sizzling; [1RB Edition]**
Colin Hunt. **Sunday Times.** London (UK): Jun 25, 2000. p. 18
Full text Abstract

6. **International Investor: Shorting the Nasdaq Could Provide Insurance For Portfolios Full of Technology Stocks --- Funds Try to Do the Opposite of the Index's Daily Performance**
By Joe Hagan. **Wall Street Journal (Europe).** Brussels: Feb 28, 2000. p. 19
Full text Abstract

7. **Advancing information retrieval through databases, fusion and information extraction**
by *McCabe, M. Catherine*, Ph.D., **George Mason University**, 2000, 138 pages; AAT 9978934
Abstract 24 Page Preview Page Image - PDF Order a copy

8. **Maxima of partial sums indexed by geometrical structures**
by *Jiang, Tiefeng*, Ph.D., **Stanford University**, 1999, 84 pages; AAT 9958121
Abstract 24 Page Preview Page Image - PDF Order a copy

9. **Factors affecting the extrusion of cereal-based foods**
by *Mathew, John Mampra*, Ph.D., **Kansas State University**, 1996, 235 pages; AAT 9720934
Abstract 24 Page Preview Page Image - PDF Order a copy

10. **Avaluacio d'indexs espectroradiometrics (entre 400 i 1100 nm) com a indicadors de la biomassa, estat hidric, contingut pigmentari i eficiencia fotosintetica de diverses especies silvestres i conreuades de l'area mediterrania** by *Filella Cubells, Iolanda*, Ph.D., **Universitat de Barcelona (Spain)**, 1995, 203 pages; AAT C541745

Abstract

1-10 of 10

Want to be notified of new results for this search? Set Up Alert

Results per page: 30

Basic Search

Tools: Search Tips Browse Topics 3 Recent Searches

(multidimensional database* or cube* or hypercube*) and (index*) and (de Search Clear

Database: Multiple databases... Select multiple databases

Date range: All dates

Limit results to: Full text documents only

Scholarly journals, including peer-reviewed About

More Search Options